

Optimizing Trial Designs for Targeted Therapies*

Thomas Ondra^{†1}, Sebastian Jobjörnsson^{†2}, Robert A. Beckman³,
Carl-Fredrik Burman^{2,4}, Franz König¹, Nigel Stallard⁵, and
Martin Posch¹

¹Center for Medical Statistics, Informatics, and Intelligent
Systems, Medical University of Vienna, Vienna, Austria

²Department of Mathematics, Chalmers University, Gothenburg,
Sweden

³Departments of Oncology and of Biostatistics, Bioinformatics,
and Biomathematics, Georgetown University Medical Center,
Washington, D.C, US

⁴Statistical Innovation, AstraZeneca R&D, Molndal, Sweden

⁵Warwick Medical School, The University of Warwick, Coventry,
UK

Abstract

An important objective in the development of targeted therapies is to identify the populations where the treatment under consideration has positive benefit risk balance. We consider pivotal clinical trials, where the efficacy of a treatment is tested in an overall population and/or in a pre-specified subpopulation. Based on a decision theoretic framework we derive optimized trial designs by maximizing utility functions. Features to be optimized include the sample size and the population in

*This project has received funding from the European Union's 7th Framework Programme for research, technological development and demonstration under the IDEAL Grant Agreement no 602552, and the InSPiRe Grant Agreement no 602144.

[†]Contributed equally.

which the trial is performed (the full population or the targeted subgroup only) as well as the underlying multiple test procedure. The approach accounts for prior knowledge of the efficacy of the drug in the considered populations using a two dimensional prior distribution. The considered utility functions account for the costs of the clinical trial as well as the expected benefit when demonstrating efficacy in the different subpopulations. We model utility functions from a sponsor's as well as from a public health perspective, reflecting actual civil interests. Examples of optimized trial designs obtained by numerical optimization are presented for both perspectives.

1. Introduction

In the development of targeted therapies the investigation of potentially predictive biomarkers is critical. If efficacy is limited to an identifiable subgroup of patients, developing a therapy for an unselected patient population is ethically problematic and will also require unnecessarily large sample sizes because of a diluted treatment effect. On the other hand, erroneously restricting a drug development program to a subpopulation is also unethical, as it excludes patients from an effective treatment. Furthermore, it will entail a financial loss for the sponsor because of unnecessary costs of biomarker development and screening and the lower prevalence of the future patient population.

Several one and two stage clinical trial designs have been proposed in which the treatment effect is tested in an overall population as well as in a subgroup of biomarker positive patients (Mandrekar and Sargent, 2009; Chen and Beckman, 2009; Freidlin et al., 2010; Mandrekar and Sargent, 2011; Freidlin et al., 2012; Ziegler et al., 2012) (see (Ondra et al., 2016) for a recent review). To account for the resulting multiple comparisons, tailored multiplicity adjustments have been developed (Song and Chi, 2007; Alosh and Huque, 2009; Bretz et al., 2009; Burman et al., 2009; Zhao et al., 2010; Spiessens and Debois, 2010; Bretz et al., 2011; Millen and Dmitrienko, 2011; Alosh and Huque, 2013). Alpha allocation has also been optimized using interim trial data and/or data external to the trial, with respect to a utility function, providing an early example of the use of decision analysis (Chen and Beckman, 2009).

In this paper we use a comprehensive decision theoretic approach to derive optimal trial designs for the development of targeted therapies. Especially, the framework allows us to assess when it is favourable to investigate the biomarker in a clinical trial and when it is actually more efficient to disregard the biomarker and to proceed with a classical trial design. This extends earlier decision theoretic methods that focused on the selection of the population for clinical trials incorporating a biomarker (Beckman et al., 2011; Krisam and Kieser, 2014; Götte et al., 2015; Kirchner et al., 2016; Krisam and Kieser, 2015; Graf et al., 2015).

Consider a setting where a single potentially predictive binary biomarker has been identified in advance, separating the full population F into biomarker positive (S) and biomarker negative (S') patients and there is prior evidence suggesting that the treatment effect may be more pronounced (or only present) in the biomarker positive group. Let λ_S and $\lambda_{S'}$, satisfying $\lambda_S + \lambda_{S'} = 1$, be the prevalences of biomarker positive and biomarker

negative patients in the full population. For this situation we consider three design options for a pivotal clinical trial: (i) The *classical design* that does not account for the biomarker status and tests for a treatment effect in the full population only, (ii) the *stratified design* that also recruits patients from the full population but where the biomarker status of each patient is determined and the treatment effect is tested in the full population and the subpopulation, and, (iii) the *enrichment design*, where patients are screened for the biomarker status and only biomarker positive patients are included in the trial.

The choice of trial design will in general not only be based on power arguments, but on the overall expected utility of different designs, accounting for the potential rewards and costs. Rewards can be quantified by the sales revenue, from a sponsor's view, or by a measure of the overall health benefit, from a public health view. The costs of the trial are determined by fixed and per patient costs as well as investments in biomarker development and the determination of the biomarker status for the patients in the trial. Based on a decision theoretic framework, we first optimize each of the three trial designs by choosing optimal sample sizes (and an optimized multiple testing procedure for the stratified design). Then, the optimal design can be selected among the three optimised designs based on their expected utilities. The optimal design choice depends on the type of utility function used (sponsor's view or public health view), the reward and cost parameters, the prior distribution on the effect sizes and the prevalence of the biomarker positive subgroup.

2. Testing Problem and Considered Trial Designs

Let $\Delta = (\delta_S, \delta_{S'})$ denote the treatment effects for the primary efficacy endpoint in the subgroup and its complement, respectively. Furthermore, let $\pi(\Delta)$ denote a prior distribution on Δ . We focus on priors that satisfy $\pi(\Delta) = 0$ for $\delta_S < \delta_{S'}$. This accounts for settings where there is some evidence that the effect size in the biomarker positive treatment group may be larger than in the biomarker negative group but not the other way around.

For simplicity, we assume that the basis of marketing authorization is a single pivotal trial. We further assume that a necessary condition for regulators to approve a drug for the populations S or F is the demonstration of a significant treatment effect in the respective population by a suitable multiple testing procedure controlling the familywise error rate (FWER) at level α in the strong sense. Consider the two null hypotheses

$$H_S : \delta_S \leq 0 \text{ and } H_F : \delta_F \leq 0,$$

where $\delta_F = \lambda_S \delta_S + \lambda_{S'} \delta_{S'}$, and let, for some trial design d , $\psi_d = (\psi_{S,d}, \psi_{F,d})$ denote a multiple testing procedure such that $\psi_{i,d} = 1$ (0) if there is a statistically significant (no significant) treatment effect in population $i = S, F$.

We consider three types of trial designs, the classical, the stratified and the enrichment design. Let $D = \{\mathbf{C}_n, \mathbf{S}_{n,\alpha_S}, \mathbf{E}_n | n \geq n_{\min}, \alpha_S \in [0, \alpha]\}$ denote the set of considered trial designs, where $\mathbf{C}_n, \mathbf{S}_{n,\alpha_S}, \mathbf{E}_n$ are defined below:

Classical design C_n refers to a classical parallel group design with per group sample size n recruiting patients from the full population and testing H_F only. H_F is tested by a non-stratified test ψ_{F,C_n} and we set $\psi_{S,C_n} = 0$.

Stratified Design S_{n,α_S} refers to a stratified design, which differs from the classical design in that analysis is stratified by the biomarker status and both hypotheses H_F and H_S are tested with a weighted multiple testing procedure with parameter α_S . As multiple testing procedure we apply the closed Spiessens-Debois' test (Song and Chi, 2007; Spiessens and Debois, 2010). This test combines the Spiessens-Debois' test for the rejection of the intersection hypothesis $H_S \cap H_F$ with the closed testing principle so as to obtain a test for the rejection of either H_S or H_F (or both). Let p_S and p_F denote unadjusted p-values for testing H_S and H_F , respectively. Here we assume that H_F is tested with a test stratified for the biomarker (in contrast to the classical design, where a non-stratified test is used as no biomarker information is available). For $\alpha_S, \alpha_F \geq 0$, the closed Spiessens-Debois' test then rejects H_S if $p_S \leq \alpha$ and either $p_S < \alpha_S$ or $p_F < \alpha_F$. Similarly, it rejects H_F if $p_F \leq \alpha$ and either $p_S < \alpha_S$ or $p_F < \alpha_F$. To control the familywise error rate at level α in the strong sense, the significance levels α_S and α_F need to satisfy

$$P_{H_S \cap H_F}(p_S \leq \alpha_S \vee p_F \leq \alpha_F) \leq \alpha. \quad (1)$$

Thus, the significance level α_F is determined by (1) if $\alpha_S \leq \alpha$ is given. Note that the corresponding function $\alpha_F(\alpha_S)$ depends on the subgroup prevalence λ_S .

We assume that in the stratified design, market authorization in the population F is not only determined by the treatment effect in F , but that regulators additionally require some evidence that there is a treatment effect in both S and S' , so that the rejection of H_F is not completely dominated by a treatment effect in a single subgroup only. Thus, we assume that the regulators' decision rule corresponds to a hypothesis test where H_F is only rejected, if, in addition, the p-values p_S and $p_{S'}$ of tests for efficacy in the two subgroups fall below corresponding thresholds τ_S and $\tau_{S'}$. The resulting modified Spiessens-Debois' test $(\psi_{S,S_{n,\alpha_S}}, \psi_{F,S_{n,\alpha_S}})$ rejects H_S if $\{p_S \leq \alpha\} \wedge \{p_S \leq \alpha_S \vee p_F \leq \alpha_F\}$ and rejects H_F if $\{p_F \leq \alpha\} \wedge \{p_S \leq \alpha_S \vee p_F \leq \alpha_F\} \wedge \{p_S \leq \tau_S \wedge p_{S'} \leq \tau_{S'}\}$. Note that this test is strictly conservative, because the consistency thresholds τ_S and $\tau_{S'}$ are not considered in the level α condition.

Enrichment Design E_n refers to an enrichment design, which differs from the classical design in that only patients from the subpopulation are recruited and only H_S is tested. In the enrichment design, H_S is tested by a test denoted by ψ_{S,E_n} and we set $\psi_{F,E_n} = 0$.

3. Utility Functions

We define utility functions that quantify the potential rewards for each of the possible trial outcomes as well as the cost of the trial. To model the rewards, we distinguish between the sponsor and the public health view, leading to different utility functions for

the two perspectives:

$$U^{(v)}(d) = \sum_{i=S, S'} \varphi_{i,d}^{(v)} - \mathbf{C}_d, \quad (2)$$

where $v = \textit{Sponsor}$ for the sponsor and $v = \textit{Public}$ for the public health view, $d \in D$ denotes the trial design, $\varphi_{i,d}^{(v)}$ the reward due to the trial outcome in subgroup $i = S, S'$ and \mathbf{C}_d the cost of the trial. The cost functions \mathbf{C}_d of the different trial designs $d \in D$ are sums of fixed costs and costs per recruited patient in the trial. Note that the per-group sample size n may vary among the three designs and below we will determine optimal sample sizes for each type of design.

For the classical design the cost function is given by

$$\mathbf{C}_{\mathbf{C}_n} = c_{\text{setup}} + 2nc_{\text{per-patient}},$$

where the setup costs of the trial c_{setup} are fixed costs and $c_{\text{per-patient}}$ are the marginal costs per patient. In the stratified design there are additional fixed costs $c_{\text{biomarker}}$ to develop the biomarker and additional per patient costs to determine the biomarker status $c_{\text{screening}}$. Thus, the cost function of the stratified design is given by

$$\mathbf{C}_{\mathbf{S}_{n, \alpha_S}} = c_{\text{setup}} + c_{\text{biomarker}} + 2n(c_{\text{per-patient}} + c_{\text{screening}}).$$

For the enrichment design the fixed costs are the same as for the stratified design. However, to recruit only biomarker positive patients one has to screen (on average) $2n/\lambda_S$ patients from the full population until $2n$ biomarker positive patients are identified. Given that the screening and determination of the biomarker status induces costs $c_{\text{screening}}$ the cost function is given by

$$\mathbf{C}_{\mathbf{E}_n} = c_{\text{setup}} + c_{\text{biomarker}} + 2n(c_{\text{per-patient}} + c_{\text{screening}}/\lambda_S).$$

3.1. The Sponsor's Utility Function

For the sponsor, the utility is the Net Present Value (NPV), which is defined as the reward (sales revenue) minus the trial costs. We model the sponsor's reward as a function of (i) the outcome of the regulatory approval process, (ii) the price the sponsor can achieve, and (iii) the size of the population the drug is licensed for.

To model (i) and (ii) we define reward functions $\varphi_i^{(\textit{Sponsor})}$ for $i = S, S'$ that specify the reward obtained in the respective population. The reward function may depend on the observed data, the design of the pivotal trial d and the prevalence of the subgroup. We model the reward as the product of the price of the drug for the treatment of a single patient times the market size. Given an overall market size N , the market sizes of the two subgroups are $\lambda_S N$ and $\lambda_{S'} N$, respectively. Furthermore, we assume that the payers are willing to pay more if a larger treatment effect was observed. If the drug is authorized for neither subgroup, both reward functions are set to zero. If the drug is authorized for the subgroup S only, the reward for the complement S' is set to zero. If the drug is authorized for the full population, we assume that the same price is charged in both subgroups.

We assume that (given that the respective hypothesis test rejects and the observed effect size exceeds a clinically relevant threshold) the price increases linearly with the observed effect size. Then the reward functions for the subgroups S and S' are

$$\begin{aligned}\varphi_{S,d}^{(Sponsor)} &= \begin{cases} \lambda_S N r_S \psi_{S,d} (\hat{\delta}_{S,d} - \mu_S)^+ & \text{if } \psi_{F,d} = 0 \\ \lambda_S N r_F \psi_{F,d} (\hat{\delta}_{F,d} - \mu_F)^+ & \text{otherwise} \end{cases} \\ \varphi_{S',d}^{(Sponsor)} &= \lambda_{S'} N r_F \psi_{F,d} (\hat{\delta}_{F,d} - \mu_F)^+, \end{aligned} \quad (3)$$

where μ_i denotes a minimal clinically relevant effect size for population $i = S, F$ and $(\cdot)^+$ denotes the positive part. $\hat{\delta}_{S,d}$ and $\hat{\delta}_{F,d}$ are the estimates of δ_S and δ_F obtained from the trial data. The constants r_i for $i = S, F$ are the marginal prices (the change in price if the observed effect size increases by one unit) and N denotes the total market size, which for the sponsor is defined as the number of future patients within the patent life of the therapy in the unselected, full population. Note that, given that efficacy is shown in the full population, a common treatment effect estimate $\hat{\delta}_F$ is used in the price function. Then the overall reward within the patent life of the therapy is given by $\varphi_{S,d}^{(Sponsor)} + \varphi_{S',d}^{(Sponsor)}$.

3.2. Public health utility function

With the public health utility function we model the utility of trial designs under the assumption that the drug is developed by public health authorities. Therefore, the utility of a trial is given by the total health benefit to the society (adjusted by the production cost of the drug) minus the cost of running the trial. We assume that the benefit of the drug is measured on a monetary scale representing the expected, accumulated (over the whole treated population) treatment effect. Costs are assumed to be the same as under the sponsor view. The reward functions for the subgroups S and S' are given by

$$\begin{aligned}\varphi_{S,d}^{(Public)} &= \begin{cases} \lambda_S N r_S \psi_{S,d} (\delta_S - \mu_S) & \text{if } \psi_{F,d} = 0 \\ \lambda_S N r_F \psi_{F,d} (\delta_F - \mu_F) & \text{otherwise} \end{cases} \\ \varphi_{S',d}^{(Public)} &= \lambda_{S'} N r_F \psi_{F,d} (\delta_F - \mu_F), \end{aligned} \quad (4)$$

The first term in the utility function (2) denotes the total benefits summed over the whole population, which are assumed to be proportional to the effect size (adjusted for a minimal relevant threshold), if the drug is authorized. The constants r_i for $i = S, F$ are the marginal benefits (the change in benefit if the effect size increases by one unit), and N denotes the size of the future (unselected) patient population. Note that the benefit depends on the actual effect sizes δ_i and not on the corresponding trial estimates $\hat{\delta}_{i,d}$, implying that the benefit may be negative if the effect size is low. A consequence of this model choice is that a public health authority will take into account the risk of false positive approvals when optimizing its trial design. Such considerations are absent when a sponsor is optimizing, since we have assumed that only the estimated effects enter its utility function.

3.3. Optimizing the expected utility

Recall that π denotes a prior distribution on the effect sizes Δ . For a given utility function $U^{(v)}$ and set of trial designs D the design optimizing the expected utility is given by

$$d^* \in \operatorname{argmax}_{d \in D} E_\pi [U^{(v)}(d)], \quad (5)$$

where

$$E_\pi[U^{(v)}(d)] = \int E_\Delta[U^{(v)}(d)] d\pi(\Delta), \quad (6)$$

Note that the expectation is first taken over the data distribution given the effect sizes Δ and then over the prior distribution π .

4. Numerical Examples

We consider parallel group designs for the comparison of means of a continuous outcome. We assume that the responses in the control and experimental treatment arms $k = C, T$ in subgroups $j = S, S'$ are normally distributed with mean $\theta_{k,j}$ and variance σ^2 . However, utilizing the central limit theorem, the model can be modified to account for many other situations. The mean treatment effects in the two subgroups are given by $\delta_j = \theta_{T,j} - \theta_{C,j}$. In the classical and the enrichment design non-stratified z-tests are performed to test H_F and H_S , respectively. In the stratified test, an elementary p-value p_S is computed from a z-test for H_S based on the observations in S and a p-value p_F from a stratified z-test for H_F stratifying by biomarker status. Then the Spiessens and Debois test is performed to adjust for multiplicity. We set $\sigma = 1$.

To be able to compute the expected utilities by numerical integration and not to have to rely on simulations, we approximated the sampling distributions for both the classical and the stratified designs by normal distributions (for the enrichment design the z-test statistic is exactly normally distributed). For the classical design, each subject recruited to the trial belongs to S with probability λ_S . Therefore, each observation in group $i = T, C$ is with probability λ_S distributed as $N(\theta_{i,S}, \sigma^2)$ and with probability $\lambda_{S'}$ distributed as $N(\theta_{i,S'}, \sigma^2)$. If the biomarker is either prognostic (such that $\theta_{C,S} \neq \theta_{C,S'}$) or predictive (such that $\delta_S \neq \delta_{S'}$) the overall treatment effect estimate $\hat{\delta}_F$ for the classical design is not exactly normal, but, for sufficiently large sample sizes, approximately normal by the central limit theorem. Because the observations are drawn from a mixture distribution, the standard deviation of $\hat{\delta}_F$ increases with the absolute differences $|\theta_{i,S} - \theta_{i,S'}|, i = T, C$. For simplicity, in the numerical examples we assume that the biomarker is predictive only but not prognostic (i.e., $\theta_{C,S} = \theta_{C,S'}$, see Appendix A for further details). For the stratified design, we assume that the subgroup estimates, $\hat{\delta}_S$ and $\hat{\delta}_{S'}$, are constructed as the sample means of exactly $\lambda_S n$ (resp. $\lambda_{S'} n$) observations per group from the subgroups S and S' . However, if patients are not selected for the trial based on biomarker status, the number of subjects from each subgroup is binomially distributed, though, for large n , the random sample sizes have only little impact and the approximation becomes accurate. Therefore, in the numerical investigations, we introduced a minimal sample size of $n_{\min} = 50$ patients per treatment arm. For the contour plot (Figure 4) optimization was performed by

	δ_S	0	δ	δ	δ
	$\delta_{S'}$	0	0	$\delta/2$	δ
"weak biomarker prior "		0.2	0.2	0.3	0.3
"strong biomarker prior"		0.2	0.6	0.1	0.1

Table 1: Prior distributions corresponding to scenarios where there is either only weak or strong prior evidence that the biomarker is predictive. The constant $\delta > 0$ parametrizes the effect sizes in the prior.

evaluating the objective function for a grid of candidate sample sizes (and α_F values for the stratified design). For the optimizations in the other plots, a further optimization step was applied by optimizing the objective functions with the R Version 3.2.4 procedure *optim* (R Core Team, 2016) using grid points as starting values.

The one-sided significance level is set to $\alpha = 0.025$ and the consistency thresholds in the multiple test for the stratified design to $\tau_S, \tau_{S'} = 0.3$. We consider discrete prior distributions $\pi_{\delta_{S,i}, \delta_{S',i}}$ on a grid $(\delta_{S,i}, \delta_{S',i}), i = 1, \dots, l$ of effect sizes and specify two priors corresponding to scenarios where there is either only weak or strong prior evidence that the biomarker is predictive. The prior distributions used in the examples are defined in Table 1 and depend on an effect size parameter δ . In the examples below we set $\delta = 0.3$ with the exception of Subsection 4.1.3 where optimal designs for other choices of δ are explored.

The reward and cost parameters in the sponsor and the public health utility function are specified via the following three cases:

Case 1 Corresponds to a large market, where the biomarker costs are negligible, i.e. $Nr_S = Nr_F = 10,000$ Million US Dollars (MUSD) per unit of efficacy and $c_{\text{screening}} = c_{\text{biomarker}} = 0$.

Case 2 Corresponds to a small market, where the biomarker costs are still negligible, i.e. $Nr_F = Nr_S = 1000$ MUSD per unit of efficacy.

Case 3 We add biomarker and screening costs, $c_{\text{screening}} = 5000$ USD per patient and $c_{\text{biomarker}} = 10$ MUSD. The reward parameters Nr_S and Nr_F are the same as in Case 2.

For all three cases we choose $c_{\text{per-patient}} = 0.05$ MUSD and $c_{\text{setup}} = 1$ MUSD. Note that the setup costs are assumed to be the same for the enrichment, classical and stratified design and therefore have no impact on the order of their expected utilities. However, they do have an impact on the sign of the utility, and thus whether any trial design is superior to no trial at all. In the reward functions (3,4) we set the minimal clinical relevant thresholds to $\mu_S = \mu_F = 0.1$, which is a third of the effect size $\delta = 0.3$ used in the prior distributions in Section 4.1.2 and 4.1.1.

4.1. Results

We discuss the optimal designs for the weak and the strong biomarker prior and the three cases specifying the cost and reward parameters.

4.1.1. Optimization under the Weak Biomarker Prior

Large market, no biomarker costs (Case 1) The optimized utilities and corresponding optimal classical, stratified and enriched designs are shown in Figure 1.

Optimal utility. For the sponsor utility function, the stratified design has the largest expected utility, except for low prevalences where the classical design is optimal. The latter is on first sight surprising, because in Case 1 we assume no biomarker costs. However, in the stratified design (in contrast to the classical design), to show efficacy in the full population, we require that p_S and $p_{S'}$ do not exceed $\tau_S = \tau_{S'} = 0.3$ (in addition to rejection of H_F in the multiple testing procedure). Thus, for low prevalences the sample size of the stratified design needs to be substantially increased to reach a sufficient power to show efficacy in F and therefore its expected utility is lower.

For the public health utility function we observe a similar pattern. However, for large λ_S the expected utility of the classical design is almost identical to that of the stratified design. This holds because the power to reject H_S in the optimized stratified design approaches the power to reject H_F in the classical design and the rewards obtained for authorization in populations S and F are similar. Why is the stratified design for the sponsor view still optimal in this case? This results from the fact that the size of the reward in the sponsor view depends on the observed rather than the true treatment effect: for trial outcomes where H_F can be rejected in the classical design but, due to the variability of estimates, $\hat{\delta}_S$ is large but $\hat{\delta}_{S'}$ is small (and thus $\hat{\delta}_F < \hat{\delta}_S$) the reward for a market authorization in S may become larger than the reward in F . However, while the classical design leads to rejection of H_F in such cases, the stratified design rejects H_S and not H_F because of the consistency threshold.

Optimal sample size. Overall, the optimized sample sizes for the public health utility function are larger than for the sponsor utility function. They are lowest for the enrichment design, and - for smaller prevalences - largest for the stratified design. For the latter, the sample size increases sharply for low prevalences. This is due to the fact that a sufficient sample size in the subgroup is required to achieve adequate power for the rejection of both H_S and H_F (for the latter due to the consistency threshold τ_S). Furthermore, the relationship of the optimal sample size and the prevalence is qualitatively different for the three designs. For both utility functions the optimal sample size is increasing in the prevalence for the enrichment design (because the gain when demonstrating efficacy in S increases), decreasing for the stratified design (because, as noted above, a sufficient sample size in S is required for the rejection of H_S and for the rejection of H_F) and non-monotone for the classical design (essentially because the effect size in population F is increasing in λ_S such that for small λ_S the expected utility does not sufficiently increase with the sample size to compensate the additional costs, while for large λ_S a smaller sample size is sufficient to achieve adequate power).

Significance levels. In the intersection hypothesis test of the optimal multiple testing




Figure 1: Weak biomarker prior and a large market with no biomarker costs (Case 1). Optimized expected utilities and sample sizes for the enrichment, classical and stratified design as functions of the prevalence for $\lambda_S \in [0.05, 0.95]$. For the stratified design, optimized levels α_S and α_F for the multiple testing procedure are given. The last row shows the overall probability (averaged over the prior) that a significant treatment effect in H_S or H_F can be shown (and, for the stratified design, that the thresholds τ_S and $\tau_{S'}$ are crossed). The priors are defined as in Table 1 with $\delta = 0.3$.




Figure 2: Weak biomarker prior and a small market with no biomarker costs (Case 2). See the legend of Figure 1.




Figure 3: Strong biomarker prior and a small market with no biomarker costs (Case 2). See the legend of Figure 1.

procedure in the stratified design α_S is larger than α_F for almost all prevalences. To make up for the lower sample sizes in the subgroup, the optimal design uses a larger α_S than α_F . For increasing prevalences, the correlation of the test statistics used to test H_S and H_F increases such that less multiplicity correction is required and both α_S and α_F increase.

Power. We define the power corresponding to a specific trial design as the overall probability (averaged over the prior) of regulatory approval in any population. This is a slight generalization of the traditional concept of power, which in the current context may be defined as the probability of regulatory approval conditional on a specific pair of subgroup effects. The power obtained by averaging over a prior has also been referred to as *assurance* (O’Hagan et al., 2005). The curves shown in Figure 1 correspond to the optimal designs. It can be seen that the power is largest for the enrichment design, followed by the stratified and the classical design and that it increases with the prevalence.

Note that for the stratified design, the probability to obtain marketing authorization in H_F is largest for intermediate values of λ_S and much lower than for the classical design if λ_S is large (even though the optimized sample sizes are similar in this case). This is due to the application of the consistency thresholds which are a more difficult to meet if one of the subgroups S or S' is small.

Small market, no biomarker costs (Case 2) Case 2 differs from Case 1 only in that the rewards Nr_F and Nr_S are reduced by a factor 10. Because of the lower rewards the optimized expected utilities are smaller compared to Case 1 (see Figure 2 for the expected utilities and optimized design parameters). They decrease even more than by a factor 10 as the trial costs are not reduced proportionally. However, the optimized sample sizes (and consequently the overall probabilities to show efficacy in the respective populations) are substantially smaller than in Case 1. Overall, the expected utilities follow a similar pattern as in Case 1 but the range of prevalences where the classical design has a higher expected utility is larger than in Case 1 for both the sponsor and the public health utility functions. The assumption of a smaller market qualitatively changes the optimized sample size of the stratified designs as a function of the prevalence. For low prevalences the optimized sample size is much lower than in Case 1: because the reward is lower, it does not pay off to invest in a large overall sample size to meet the threshold τ_S in the subpopulation. This is also reflected in the optimized significance levels α_S and α_F , which give more weight to H_F than in Case 1.

Small market with biomarker costs (Case 3) Note that the addition of biomarker costs has no impact on the expected utility of the optimal classical design (as it does not require the biomarker). However, the expected utilities of the enrichment and the stratified design become smaller compared to Case 2 because of the additional costs. Therefore, the classical design now dominates the stratified design for a broader range of (small) values of λ_S and the stratified design becomes optimal only for larger values of λ_S (see Figure S1 in the Supplementary Material). In the public health view, the classical design dominates the stratified design also for very large values of λ_S : even though the classical design leads to lower expected rewards compared to the stratified design (since

the latter is more likely to lead to market authorization for too large a population), this is compensated by the lower costs because no biomarker is required. In the sponsor view in contrast, the difference between the expected rewards of the stratified and the classical design is larger because it is determined by observed treatment effects (see also the discussion of expected utilities in Case 1, where a similar pattern is observed). Therefore, the stratified design dominates also for large values of λ_S .

Moreover, the biomarker costs lead to a reduction in sample size compared to Case 2.

4.1.2. Optimization under the Strong Biomarker Prior

First, note that the expected utility and optimal sample size of the enrichment design is the same for the weak and the strong biomarker prior because the prior distribution on the treatment effect in S is identical in both.

Large market, no biomarker costs (Case 1) For the sponsor's utility function the stratified design is still optimal, with the exception of very low prevalences (see Figure S2 in the Supplementary Material). In contrast, for the public health utility function, the enrichment and the stratified design have almost identical expected utilities unless the prevalence is small.

Small market, no biomarker costs (Case 2) While, as in Case 1, the stratified design is optimal for the sponsor view for all but very low prevalences, the difference between the expected utilities of the stratified and enrichment design is small.

In contrast, for the public health utility function the enrichment design achieves the highest expected utility (see Figure 3). Furthermore, for very low prevalences, none of the trial designs has a positive expected utility in the public health view and the optimal strategy is to perform no trial at all. For the sponsor view it is still optimal to perform a trial in the unselected population, albeit with the minimal sample size if the prevalence is small. This is due to the assumption that the NPV depends on the observed effect sizes, which implies that the sponsor benefits from a high variability of the treatment effect estimates.

Note that the optimal test in the stratified design gives most weight on H_F for low and on H_S for large prevalences. This holds for both the public health and the sponsor utility function.

Small market, biomarker costs (Case 3) The pattern is very similar to Case 2, however, the range of λ_S values where the classical (for the sponsor utility) or no trial (for the public health utility) are optimal becomes larger (see Figure S3 in the Supplementary Material).

4.1.3. Optimized Designs for Varying Effect Sizes

Figure 4 shows the optimal design as function of the prevalence λ_S and the effect size parameter δ which parametrizes the effect sizes in the weak and strong biomarker priors in Table 1. Under the sponsor view, either the classical or the stratified design is optimal




Figure 4: Optimal designs for different combinations of the prevalence $\lambda_S \in [0.05, 0.95]$ and effect size parameter $\delta \in [0, 1]$. Optimized utilities for the sponsor and the public health authority are shown for both the weak and the strong biomarker prior (as defined in Table 1) under the three different cost structures defined by Cases 1, 2 and 3. The colour in a specific point indicate the type of the optimal design. Grey areas correspond to regions where all optimized designs have negative utilities, implying that the optimal choice is to perform no trial.

while the enrichment design never maximizes the expected utility. Surprisingly, even for $\delta = 0$, it is never optimal from a sponsor point of view to conduct no study at all in the scenarios investigated. This is due to the fact that a false positive, even though unlikely, may lead to a large reward. Therefore the optimal sample size is the minimal sample size n_{\min} in these scenarios. This choice minimizes the costs and maximizes the variability of estimates.

For the public health view in contrast, for very low effect sizes, the optimal decision is to perform no trial at all. Under the weak biomarker prior, the enrichment design is optimal under the public health view only in the scenarios without biomarker costs, for small δ and large enough prevalences (such that the population that will benefit from a new treatment in the future is large enough). For larger effect sizes the classical design is optimal for very low and very large prevalences and the stratified design otherwise.

Under the strong biomarker prior and intermediate δ , the public health utility is optimized by the enrichment design (unless the prevalence is too low and the classical design dominates). For larger δ the stratified design is optimal, again with the exception of very low prevalences. In addition, in the scenario with biomarker costs the classical design becomes optimal for large prevalences.

5. Discussion

The current study suggests decision-theoretic models for optimizing confirmatory biomarker trials, both from a sponsor and a public health perspective. Furthermore, it explores the potential discrepancies between the two perspectives.

The optimized designs depend sensitively on the particular configuration of parameter values. Besides the priors on the effect sizes, the assumptions on the market size and costs have a substantial impact on the optimized designs. Therefore, formulating simple rules of thumb for trial designs is hardly feasible. However, a few general observations can be made. The optimized sample sizes for the public health utility function are consistently larger than for the sponsor utility (assuming the same costs, market size and reward parameters r_F, r_S in both utility functions). This finding is likely due to the fact that sponsor benefit is based on the estimate of the benefit in the trial, whereas the public health benefit depends on the actual benefit. Thus the public health perspective implies a higher standard for the evidence. This finding provides a quantitative basis for the qualitative observation that health authorities tend to require a higher standard of evidence than desired by some sponsors. Mechanism design theory could potentially be applied to try to create mechanisms which align the incentives more completely.

Furthermore, for very low prevalences, the classical design outperforms the designs that are based on the biomarker. However, in these scenarios the expected utility for all designs can be negative in the public health perspective and so weakly positive in the sponsor perspective that the sponsor would allocate its resources elsewhere as well.

We find that in the sponsor view the enrichment designs never maximize the NPV in the considered scenarios. This is due to the fact that the sponsor may benefit from an authorization in the larger population even if the treatment is effective in the subpopulation only. For similar reasons, even under the global null hypothesis the strategy to

perform a trial (with minimum sample size) gives a positive NPV in the sponsor view (a phenomenon that was observed also in other contexts (Posch and Bauer, 2013)).

In the public health view the enrichment designs are optimal for a range of scenarios. Especially, if there is sufficient prior medical understanding that the biomarker negative subpopulation is unlikely to be positively affected by the drug, it can be a waste of resources to conduct the trial in this population. Ethical considerations reinforce this, as it can be argued that genuine informed consent (Burman and Carlberg, 2010) implies that patients should not be randomised if their expected utility is higher on standard of care than on randomised trial medication. On the other hand, in particular when subpopulations can be expected to be similar in efficacy, it is not always worthwhile to conduct biomarker screening. In fact, there is an increased risk in a stratified trial that the treatment is rejected in the biomarker negative subpopulation due to chance. Still, in situations with genuine uncertainty about the relative efficacy in the two subpopulations, biomarker determination and stratified designs may have a large value. An obvious extension of our model is to allow for trial adaptations, potentially closing the biomarker negative part of the trial at an interim, in case results are negative (Brannath et al., 2009; Bauer et al., 2016).

When applying the presented framework to practical design decisions, the different model components should be scrutinized. In the numerical example we have assumed for simplicity that the biomarker is not prognostic but in practice this will often not be the case. If the biomarker is also prognostic, the variability of the effect size estimates will be increased with a consequent decrease in the expected utility of the classical design.

As regards the market size for the sponsor, N denotes the number of patients, determined by the patent life, for which full payment will be received upon regulatory approval. On the other hand, for the public health authority, N denotes the total number of future patients. In an extended model, N could be fixed to always be the total number of future patients and a factor could be added next to N for the sponsor. This factor would then represent the fraction of patients corresponding to the patent life, and could be made to depend on the choice of trial design in various ways. For example, in the enrichment design we accounted only for the screening costs arising from the determination of the biomarker status of patients. However, if the restriction of the trial population leads to slower recruitment and consequently a later authorization of the drug, the result will be a reduction of the market size and the remaining patent life. This, in turn, may reduce the potential reward in different ways for the two perspectives. Another simplification made in our framework is the assumption of a zero discount rate for the sponsor. In practice, a commercial sponsor would use a non-zero rate to discount future revenues, which would lead to a further reduction of its expected utility as compared to a public health authority.

In the considered model we assumed that the subgroup prevalences in the trial are the same as in the total patient population. However, unless the recruitment is stratified by subgroup, the actual prevalence in the trial will be stochastic. Furthermore, the propensity to participate in the trial may vary between subpopulations. While our results are generally robust to random variations in the prevalence, varying propensities for trial participation may lead to a biased estimation of the effect size in the full population

in the classical design (and also the stratified design, if an overall effect estimate based on observed prevalences is computed). The question of generalizability of trial results to general patient populations is however not specific to the development of targeted therapies but a more general issue.

We did not explicitly incorporate a benefit risk evaluation of the treatment into the model. However, the parameters μ_S and μ_F in the reward functions (3,4) can be interpreted as the minimal treatment effects that compensate the "costs" of the treatment, as the burden of treatment, side effects and monetary costs. While these are considered as given in our model, they could alternatively be estimated from clinical trial data.

We modelled the sponsor and public health utility as essentially linear functions of the observed and true effect size, respectively. From a commercial perspective this can be reasonable for scenarios where no alternative treatment options are available. However, if competitor products are on the market, the model may need to be modified because the market share, in terms of number of doses prescribed, and not only the price or benefit per patient may depend on the effect size. This can be incorporated by models where the market share is a function of the posterior distribution of efficacy (and possibly safety) parameters (Gittins and Pezeshk, 2000; Kikuchi et al., 2008). Another aspect of our model for reimbursement concerns the pricing. Although NICE in the UK indicates that they, in our situation, would accept a price proportional to net benefit, payers in other countries may use other price models, possibly closer to a constant price. As an alternative to our linear sales model, an aggregated commercial model could be plugged in and similar optimization could be performed.

Finally, we note that the case of a very low prevalence, small market size and no biomarker costs mimics the situation of a rare disease, except there is no complementary subgroup S' . Therefore, our results could be seen to suggest that the investigation of rare diseases is not recommended in either perspective. Consequently, the question arises if research in rare diseases should receive special priority and be subsidised by society such that drug development occurs even though the expected utility to society is negative, or in some case weakly positive but less positive than other alternative expenditures. However, this argument raises ethical questions because the purely utilitarian viewpoint that underlies the decision theoretic framework does not account for other ethical principles as fairness and justice. Similar issues arise for small subgroups of common diseases, an increasing issue in cancer given the fact it being subdivided into many small molecular subclasses. In the case of cancer, increased benefit due to matching between molecular subgroups and targeted therapies may mollify this issue, but this remains to be seen in individual cases, so that the ethical and public policy dilemma may still be present.

References

- Alosh, M. and M. F. Huque (2009). A flexible strategy for testing subgroups and overall population. *Statistics in Medicine* 28(1), 3–23.
- Alosh, M. and M. F. Huque (2013). Multiplicity considerations for subgroup analysis subject to consistency constraint. *Biometrical Journal* 55(3), 444–462.

- Bauer, P., F. Bretz, V. Dragalin, F. König, and G. Wassmer (2016). Twenty-five years of confirmatory adaptive designs: opportunities and pitfalls. *Statistics in Medicine* 35, 325–347.
- Beckman, R. A., J. Clark, and C. Chen (2011). Integrating predictive biomarkers and classifiers into oncology clinical development programmes. *Nature Reviews Drug Discovery* 10(10), 735–748.
- Brannath, W., E. Zuber, M. Branson, F. Bretz, P. Gallo, M. Posch, and A. Racine-Poon (2009). Confirmatory adaptive designs with Bayesian decision tools for a targeted therapy in oncology. *Statistics in Medicine* 28, 1445–1463.
- Bretz, F., W. Maurer, W. Brannath, and M. Posch (2009). A graphical approach to sequentially rejective multiple test procedures. *Statistics in Medicine* 28(4), 586–604.
- Bretz, F., M. Posch, E. Glimm, F. Klinglmueller, W. Maurer, and K. Rohmeyer (2011). Graphical approaches for multiple comparison procedures using weighted Bonferroni, Simes, or parametric tests. *Biometrical Journal* 53(6), 894–913.
- Burman, C., C. Sonesson, and O. Guilbaud (2009). A recycling framework for the construction of Bonferroni-based multiple tests. *Statistics in Medicine* 28(4), 739–761.
- Burman, C.-F. and A. Carlberg (2010). Future challenges in design and ethics of clinical trials. In *Pharmaceutical Sciences Encyclopedia*, Volume 51, pp. 1–28.
- Chen, C. and R. A. Beckman (2009). Hypothesis testing in a confirmatory phase iii trial with a possible subset effect. *Statistics in Biopharmaceutical Research* 1, 431–440.
- Freidlin, B., L. M. McShane, and E. L. Korn (2010). Randomized clinical trials with biomarkers: Design issues. *JNCI Journal of the National Cancer Institute* 102(3), 152–160.
- Freidlin, B., L. M. McShane, M.-Y. C. Polley, and E. L. Korn (2012). Randomized phase II trial designs with biomarkers. *Journal of Clinical Oncology* 30(26), 3304–3309.
- Gittins, J. and H. Pezeshk (2000). A behavioral Bayes method for determining the size of a clinical trial. *Drug Information Journal* 34, 355–363.
- Graf, A. C., M. Posch, and F. Koenig (2015). Adaptive designs for subpopulation analysis optimizing utility functions. *Biometrical Journal* 57, 76–89.
- Götte, H., M. Donica, and G. Mordenti (2015). Improving probabilities of correct interim decision in population enrichment designs. *Journal of Biopharmaceutical Statistics* 25(5), 1020–38.
- Kikuchi, T., H. Pezeshk, and J. Gittins (2008). A Bayesian cost-benefit approach to the determination of sample size in clinical trials. *Statistics in Medicine* 27(1).
- Kirchner, M., M. Kieser, H. Götte, and A. Schöler (2016). Utility-based optimization of phase ii/iii programs. *Statistics in Medicine* 35(2).

- Krisam, J. and M. Kieser (2014). Decision rules for subgroup selection based on a predictive biomarker. *Journal of Biopharmaceutical Statistics* 24(1), 188–202.
- Krisam, J. and M. Kieser (2015). Optimal decision rules for biomarker-based subgroup selection for a targeted therapy in oncology. *Int. J. Mol. Sci.* 16(5), 10354–75.
- Mandrekar, S. J. and D. J. Sargent (2009). Clinical trial designs for predictive biomarker validation: One size does not fit all. *Journal of Biopharmaceutical Statistics* 19(3), 530–542.
- Mandrekar, S. J. and D. J. Sargent (2011). Design of clinical trials for biomarker research in oncology. *Clinical Investigation* 1(12), 1627–1636.
- Millen, B. A. and A. Dmitrienko (2011). Chain procedures: A class of flexible closed testing procedures with clinical trial applications. *Statistics in Biopharmaceutical Research* 3(1), 14–30.
- O’Hagan, A., J. W. Stevens, and M. J. Campbell (2005). Assurance in clinical trial design. *Pharmaceutical Statistics* 4, 187–201.
- Ondra, T., A. Dmitrienko, T. Friede, A. Graf, F. Miller, N. Stallard, and M. Posch (2016). Methods for identification and confirmation of targeted subgroups in clinical trials: a systematic review. *Journal of Biopharmaceutical Statistics* 26(1), 99–119.
- Posch, M. and P. Bauer (2013). Adaptive budgets in clinical trials. *Statistics in Biopharmaceutical Research* 5, 282–292.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Song, Y. and G. Y. H. Chi (2007). A method for testing a prespecified subgroup in clinical trials. *Statistics in Medicine* 26(19), 3535–3549.
- Spiessens, B. and M. Debois (2010). Adjusted significance levels for subgroup analyses in clinical trials. *Contemporary Clinical Trials* 31(6), 647–656.
- Zhao, Y. D., A. Dmitrienko, and R. Tamura (2010). Design and analysis considerations in clinical trials with a sensitive subpopulation. *Statistics in Biopharmaceutical Research* 2(1), 72–83.
- Ziegler, A., A. Koch, K. Krockenberger, and A. Großhennig (2012). Personalized medicine using DNA biomarkers: a review. *Human Genetics* 131(10), 1627–1638.

Appendix

Appendix A. Computation of Expected Utilities

We derive the expected utilities for a given effect size Δ for the enrichment, the classical and the stratified design. The overall expected utilities are then obtained from (6) by integrating over the prior distribution.

Enrichment Design For the enrichment design, $\psi_{F, \mathbf{E}_n} = \varphi_{S', \mathbf{E}_n}^{(Sponsor)} = 0$. Thus, the utility is given by

$$U^{(Sponsor)}(\mathbf{E}_n) = \varphi_{S, \mathbf{E}_n}^{(Sponsor)} - \mathbf{C}_{\mathbf{E}_n} = \lambda_S N r_S \psi_{S, \mathbf{E}_n} (\hat{\delta}_{S, \mathbf{E}_n} - \mu_S)^+ - \mathbf{C}_{\mathbf{E}_n}.$$

Integrating over the resulting truncated normal distribution, the expected utility given Δ is given by

$$E_{\Delta}[U^{(Sponsor)}(\mathbf{E}_n)] = \lambda_S N r_S \left((1 - \Phi(\kappa))(\delta_S - \mu_S) + V[\hat{\delta}_{S, \mathbf{E}_n}]^{1/2} \phi(\kappa) \right) - \mathbf{C}_{\mathbf{E}_n},$$

where ϕ denotes the density and Φ denotes the cumulative distribution function of the standard normal distribution, $V[\hat{\delta}_{S, \mathbf{E}_n}] = 2\sigma^2/n$ and

$$\kappa = \left[\max \left(\Phi^{-1}(1 - \alpha) V[\hat{\delta}_{S, \mathbf{E}_n}]^{1/2}, \mu_S \right) - \delta_S \right] / V[\hat{\delta}_{S, \mathbf{E}_n}]^{1/2}.$$

Similarly, for the public health view utility function we obtain

$$E_{\Delta}[U^{(Public)}(\mathbf{E}_n)] = \lambda_S N r_S (\delta_S - \mu_S) \left(1 - \Phi \left(\Phi^{-1}(1 - \alpha) - \delta_S V[\hat{\delta}_{S, \mathbf{E}_n}]^{-1/2} \right) \right) - \mathbf{C}_{\mathbf{E}_n}.$$

Classical Design In the classical design, $\psi_{S, \mathbf{C}_n} = 0$, $\varphi_{S, \mathbf{C}_n}^{(Sponsor)} = \varphi_{S', \mathbf{C}_n}^{(Sponsor)}$ and

$$U^{(Sponsor)}(\mathbf{C}_n) = \varphi_{S, \mathbf{C}_n}^{(Sponsor)} + \varphi_{S', \mathbf{C}_n}^{(Sponsor)} - \mathbf{C}_{\mathbf{C}_n} = N r_F \psi_{F, \mathbf{C}_n} (\hat{\delta}_{F, \mathbf{C}_n} - \mu_F)^+ - \mathbf{C}_{\mathbf{C}_n}$$

If the mean response in S and S' differ, it follows that the observations in the experimental treatment and control group follow a mixture distribution of two normal distributions. Therefore, the variance of $\hat{\delta}_{F, \mathbf{C}_n}$ in the classical design is given by

$$V[\hat{\delta}_{F, \mathbf{C}_n}] = (2\sigma^2 + \lambda_S \lambda_{S'} ((\theta_{T, S} - \theta_{T, S'})^2 + (\theta_{C, S} - \theta_{C, S'})^2)) / n.$$

Thus, the expected utility given effect sizes Δ for the classical design is given by

$$E_{\Delta}[U^{(Sponsor)}(\mathbf{C}_n)] = N r_F \left((1 - \Phi(\kappa))(\delta_F - \mu_F) + V[\hat{\delta}_{F, \mathbf{C}_n}]^{1/2} \phi(\kappa) \right) - \mathbf{C}_{\mathbf{C}_n},$$

where $\delta_F = \lambda_S \delta_S + \lambda_{S'} \delta_{S'}$ and $\kappa = \left(\max(\Phi^{-1}(1 - \alpha) V[\hat{\delta}_{F, \mathbf{C}_n}]^{1/2}, \mu_F) - \delta_F \right) / V[\hat{\delta}_{F, \mathbf{C}_n}]^{1/2}$. Similarly, for the public health utility function,

$$E_{\Delta}[U^{(Public)}(\mathbf{C}_n)] = N r_F (\delta_F - \mu_F) \left(1 - \Phi \left(\Phi^{-1}(1 - \alpha) - \delta_F V[\hat{\delta}_{F, \mathbf{C}_n}]^{-1/2} \right) \right) - \mathbf{C}_{\mathbf{C}_n}.$$

Stratified Design The utility of the stratified design is given by

$$U^{(Sponsor)}(\mathbf{S}_{n,\alpha_S}) = -C_{\mathbf{S}_{n,\alpha_S}} + \begin{cases} \lambda_S N r_S \psi_{S,\mathbf{S}_{n,\alpha_S}} (\hat{\delta}_{S,\mathbf{S}_{n,\alpha_S}} - \mu_S)^+ & \text{if } \psi_{F,\mathbf{S}_{n,\alpha_S}} = 0 \\ N r_F \psi_{F,\mathbf{S}_{n,\alpha_S}} (\hat{\delta}_{F,\mathbf{S}_{n,\alpha_S}} - \mu_F)^+ & \text{otherwise.} \end{cases}$$

The utility of the stratified design depends on the stratified treatment effect estimate in the full population (in the following we shorten the notation by dropping the design index, $\hat{\delta}_F := \hat{\delta}_{F,\mathbf{S}_{n,\alpha_S}}, \hat{\delta}_S := \hat{\delta}_{S,\mathbf{S}_{n,\alpha_S}}, \hat{\delta}_{S'} := \hat{\delta}_{S',\mathbf{S}_{n,\alpha_S}}, \psi_F := \psi_{F,\mathbf{S}_{n,\alpha_S}}, \psi_S := \psi_{S,\mathbf{S}_{n,\alpha_S}}$) which is a weighted sum of $\hat{\delta}_S$ and $\hat{\delta}_{S'}$. The expected utility given the effect sizes Δ is given by

$$E_{\Delta}[U^{(Sponsor)}(\mathbf{S}_{n,\alpha_S})] = N r_F E_{\Delta} [\psi_F (\hat{\delta}_F - \mu_F)^+] + \lambda_S N r_S E_{\Delta} [(1 - \psi_F) \psi_S (\hat{\delta}_S - \mu_S)^+] - C_{\mathbf{S}_{n,\alpha_S}}.$$

and can be computed by numeric integration: Let $A_F(n, \alpha_S; \sigma, \alpha, \lambda_S, \tau_S, \tau_{S'}, \mu_F)$ denote the region in \mathbb{R}^2 where $\psi_F(\hat{\delta}_F(Z_S, Z_{S'}) - \mu_F)^+ > 0$ and let $A_S(n, \alpha_S; \sigma, \alpha, \lambda_S, \tau_S, \tau_{S'}, \mu_S)$ be the region where $(1 - \psi_F)\psi_S(\hat{\delta}_S(Z_S) - \mu_S)^+ > 0$, where $\hat{\delta}_F(Z_S, Z_{S'})$ is the stratified treatment effect estimate and $Z_S, Z_{S'}$ the z-statistics computed from the observations in S and S' respectively. Then

$$E_{\Delta} [\psi_F (\hat{\delta}_F(z_S, z_{S'}) - \mu_F)^+] = \iint_{A_F} (\hat{\delta}_F(z_S, z_{S'}) - \mu_F) \phi(z_S) \phi(z_{S'}) dz_S dz_{S'},$$

and

$$E_{\Delta} [(1 - \psi_F) \psi_S (\hat{\delta}_S(Z_S) - \mu_S)^+] = \iint_{A_S} (\hat{\delta}_S(z_S) - \mu_S) \phi(z_S) \phi(z_{S'}) dz_S dz_{S'}.$$

The shapes of the regions A_F and A_S depend on the specific values of the parameters and the design variables ($\alpha_S, \alpha_F, \tau_S, \tau_{S'}$ and n). However, the regions may in all cases be described by means of a finite number of straight lines, implying that the expected values above can be computed using standard software for numerical quadrature in two dimensions. But since the integrands are linear in $z_{S'}$ and $Z_{S'}$ follows a normal distribution, one-dimensional integration may be carried out analytically in the $z_{S'}$ -direction before applying a numerical method. This leads to faster numerical evaluations, which is useful when investigating how the optimal solution changes over the parameter space.

For the public health view the expected utility given Δ may be written as

$$E_{\Delta}[U^{(Public)}(\mathbf{S}_{n,\alpha_S})] = N r_F (\delta_F - \mu_F) E_{\Delta} [\psi_F] + \lambda_S N r_S (\delta_S - \mu_S) E_{\Delta} [(1 - \psi_F) \psi_S] - C_{\mathbf{S}_{n,\alpha_S}}.$$

The numerical evaluation is similar to the evaluation of the conditional expectation of the utility of the stratified design for the sponsor's view.

**Supplementary Material for
Optimizing Trial Designs for Targeted Therapies**

Thomas Ondra, Sebastian Jobjörnsson, Robert A. Beckman, Carl-Fredrik
Burman, Franz König, Nigel Stallard, and Martin Posch




Figure S1: Weak biomarker prior and a small market with biomarker costs (Case 3). Optimized expected utilities and sample sizes for the enrichment, classical and stratified design as functions of the prevalence for $\lambda_S \in [0.05, 0.95]$. For the stratified design, optimized levels α_S and α_F for the multiple testing procedure are given. The last row shows the overall probability (averaged over the prior) that a significant treatment effect in H_S or H_F can be shown (and, for the stratified design, that the thresholds τ_S and $\tau_{S'}$ are crossed).




Figure S2: Strong biomarker prior and a large market with no biomarker costs (Case 1). See the legend of Figure S1.




Figure S3: Strong biomarker prior and a small market with biomarker costs (Case 3). See the legend of Figure S1.